

WHAT IS CLAIMED IS:

1. A dynamically scalable storage system, comprising:
 - a first plurality of disk drives for storing distributed parity groups, each distributed parity group comprising storage blocks, said storage blocks comprising one or more data blocks and a parity block, each of said storage blocks stored on a separate disk drive such that no two storage blocks from a given distributed parity group reside on the same disk drive;
 - 5 file system metadata to describe a location of each of said storage blocks by specifying a disk identifier and a logical block identifier of a first logical block of each storage block;
 - 10 a processor configured to manage said file system metadata, said processor configured to recognize when a new disk drive is added to said first plurality of disk drives to produce a second plurality of disk drives, said processor further configured to cause one or more existing storage blocks from said first plurality of disk drives to be moved to said new disk drive to balance usage of all drives in said second plurality of disk drives, said processor further configured to update said file system metadata to point to storage blocks moved to said new disk drive.
- 15 2. The dynamically scalable storage system of Claim 1, said processor further configured to allocate space for at least a portion of newly created distributed parity groups on said new drive.
- 20 3. The dynamically scalable storage system of Claim 1, said processor further configured to use any free logical blocks on said new disk drive to store one or more storage blocks.
- 25 4. The dynamically scalable storage system of Claim 3, wherein said new disk drive is larger in capacity than any drive in said first plurality of disk drives.
- 30 5. The dynamically scalable storage system of Claim 1, wherein a size of a first distributed parity group is larger than a size of a second distributed parity group within a first file.
6. The dynamically scalable storage system of Claim 1, wherein said new disk drive is provided to a Fibre Channel network.

7. The dynamically scalable storage system of Claim 1, wherein a file is organized as one or more distributed parity groups of varying size.

8. The dynamically scalable storage system of Claim 1, wherein a file is organized as one or more distributed parity groups having different extents.

5 9. The dynamically scalable storage system of Claim 1, wherein an extent of a first distributed parity group of a file is larger than an extent of a second distributed parity group of said file.

10 10. A method for managing a storage array, comprising:

recognizing a new storage device;

10 adding said new storage device to a list of previously-available storage devices to produce a list of currently-available storage devices;

determining a size of a new parity group, said size describing a number of data blocks in said new parity group;

15 computing a parity block for said parity group;

storing one of said data blocks of said parity group on said new storage device;

combining a first distributed parity group with a second distributed parity group to produce at least one new parity group; and

20 storing one data block of said at least one parity group on said new storage device.

11. The method of Claim 10, further comprising storing metadata to describe a disk and logical block location of each of said data blocks and said parity block.

12. The method of Claim 10, wherein said new storage device comprises a new disk drive.

25 13. The method of Claim 10, wherein said new disk drive is larger in capacity than other disk drives in said list of previously-available storage devices, and wherein said one data block of said at least one parity group can be stored on any unused logical block of said new disk drive.

14. The method of Claim 10, further comprising: recognizing that a selected 30 storage device has gone offline, removing said selected storage device from said list of currently-available storage devices to produce a list of remaining storage devices;

reconstructing data stored on said selected storage device; storing said reconstructed data on one or more of said remaining storage devices; and updating file system metadata to facilitate locating said reconstructed data.